

Syllabus subject to change

CS 411: Database Systems

Course Description

As the first introductory course for databases, this course studies the fundamentals of using and implementing relational and non-relational database management systems. First, from the user perspective (i.e., how to use a database system), the course will discuss conceptual data modeling, physical data modeling, computing on data, designing schemas, querying databases, and manipulating databases. We will study both the relational and “NoSQL” databases such as document and graph databases. Further, from the system perspective (i.e., how to design and implement a database system), the course will study data representation, data accessing and indexing, query processing and optimization, and transaction processing.

Course Prerequisite

- Background: CS225 Data Structures.
- Programming: You should be comfortable in programming and can quickly learn a new language if needed. You should be able to use major programming languages such as Python and Java.

Course Goals

By the end of the course, you will be able to:

- Use a database system for storing, managing, and querying data.
- Design and create a data-driven application using a database system.
- Identify the differences between relational and non-relational (NoSQL) systems and their target usage and application settings.
- Describe the internal architecture and key components of a database system.
- Design and implement data indexing, query processing and optimization, and transaction processing algorithms.

Textbook

[Database Systems: The Complete Book, 2nd edition](#), by Hector Garcia-Molina, Jeffrey D. Ullman, and Jennifer D. Widom.

Course Schedule

Week	Duration (MM/DD – MM/DD)	Topic	Relevant Concepts and Techniques
1	1/15 – 1/21	Introduction	Course format, exams, projects, database introduction and overview.
2	1/22 – 1/28	Conceptual Data Modeling	Data modeling, ER model, UML.
3	1/29 – 2/4	Physical Data Modeling	Relational model, translating from ER to relational model, object-oriented and object-relational models, NoSQL models-- key-value model, document mode, graph model.
4	2/5 – 2/11	Computing on Data	Relational algebra, map-reduce framework.
5	2/12 – 2/18	Designing Schemas	Schema redundancy and anomalies, functional dependencies, normal forms, Boyce-Cod Normal Form, Third Normal Form, lossless and dependency-preserving decomposition, multi-valued dependencies.
6	2/19 – 2/25	Querying Databases: The Relational Way	SQL, basic select-from-where, null values, subqueries, set operations, dealing with duplicates, aggregates.
7	2/26 – 3/4	Querying Databases: Non-relational Way	Querying document databases, MongoDB querying, basic operations, handling complex structures, aggregate framework, the case for joins.
8	3/5 – 3/11	Querying Databases: Non-relational Way	Querying graph databases, Neo4j querying, basic operations, advanced capabilities.
9	3/12 – 3/18	Manipulating Databases	Creating and modifying databases, views, constraints, checks and assertions, triggers.
10	3/19 – 3/25	Spring Break	
11	3/26 – 4/1	Accessing and Indexing Data	Accessing data, indexing data, index-sequential access method, tree-based indexing, binary tree, b+tree, quad tree, hash-based indexing, linear hashing, extensible hashing.
12	4/2 – 4/8	Query Processing	Database machine cost models, query operators, one-pass algorithms.
13	4/9 – 4/15	Query Processing	Two-pass algorithms, index-based algorithms.
14	4/16 – 4/22	Query Optimization	Algebraic laws, optimizer, rule-based optimization, cost-based optimization, dynamic programming, physical query plans, estimating sizes.
15	4/23 – 4/29	Transaction Management	Transactions, correctness of transactions, recovery, undo

			logging, redo logging, undo/redo logging.
16	4/30 – 5/6	Conclusion	Review and discussion of the state of industry and research.

Elements of This Course

The course is comprised of the following elements:

- Lecture Video. Each topic will be covered through about 6 to 10 individual lecture videos that range from 5 to 20 minutes long each.
- Assignments. Assignments are given in the form of practice quiz. Each week has one practice quiz with 10 questions. Practice quizzes must be attempted and will be graded based on completion (you will receive full credit as long as you make a meaningful effort taking these quizzes). You will have unlimited attempts on practice quiz.
- Quizzes. Each week has one graded quiz with 10 questions. Graded quizzes contribute to the course overall grading.
- Programming Assignments. Four programming assignments are given in this class. Programming assignments will be autograded.
- Exams. Two proctored exams (via ProctorU) will be given in this class.
- Project. A semester-long project, with deliverables scheduled throughout the semester.

Assignment Deadlines

For all assignment deadlines, please refer to the Course Deadlines, Late Policy, and Academic Calendar page.

Grading Distribution and Scale

Grading Distribution

Your final grade will be calculated based on the activities listed in the table below. Your official final course grade will be listed in [Enterprise](#). The course grade you see displayed in Coursera may not match your official final course grade.

Component	Percentage
Assignment	5%
Graded Quizzes	25%
Programming Assignments	20%
Proctored Exams	25%
Project	25%
Total	100%

Grading Scale

Total	Grade
90 - 100	A (A-, A, A+)
80 - 89	B (B-, B, B+)
70 - 79	C (C-, C, C+)
60 - 69	D (D-, D, D+)
Below 60	F

Student Code and Policies

A student at the University of Illinois at the Urbana-Champaign campus is a member of a University community of which all members have at least the rights and responsibilities common to all citizens, free from institutional censorship; affiliation with the University as a student does not diminish the rights or responsibilities held by a student or any other community member as a citizen of larger communities of the state, the nation, and the world. See the [University of Illinois Student Code](#) for more information.

Academic Integrity

All students are expected to abide by [the campus regulations on academic integrity found in the Student Code of Conduct](#). These standards will be enforced and infractions of these rules will not be tolerated in this course. Sharing, copying, or providing any part of a homework solution or code is an infraction of the University's rules on academic integrity. We will be actively looking for violations of this policy in homework and project submissions. Any violation will be punished as severely as possible with sanctions and penalties typically ranging from a failing grade on this assignment up to a failing grade in the course, including a letter of the offending infraction kept in the student's permanent university record.

Again, a good rule of thumb: *Keep every typed word and piece of code your own*. If you think you are operating in a gray area, you probably are. If you would like clarification on specifics, please contact the course staff.

Disability Accommodations

Students with learning, physical, or other disabilities requiring assistance should contact the instructor as soon as possible. If you're unsure if this applies to you or think it may, please contact the instructor and [Disability Resources and Educational Services \(DRES\)](#) as soon as possible. You can contact DRES at 1207 S. Oak Street, Champaign, via phone at (217) 333-1970, or via email at disability@illinois.edu.